

Punishment catalyzes the evolution of cooperation

Arend Hintze^{1,3} and Christoph Adami^{1,2,3,*}

¹Department of Microbiology and Molecular Genetics

²Department of Physics and Astronomy

³BEACON Center for the Study of Evolution in Action

Michigan State University, East Lansing, MI

* E-mail: adami@msu.edu

Abstract

The evolution of cooperation has been a perennial problem in evolutionary biology because cooperation can be undermined by selfish cheaters (or “free riders”) that profit from cooperators but do not invest any resources themselves. Evolutionary game theory has been able to show that under certain conditions, cooperation nonetheless evolves stably, for example if players have the opportunity to punish free riders that benefit from a public good, but refuse to pay into a common pool. In the public goods game, cooperation emerges naturally if the synergy of the public good is sufficiently high. However, a necessary high synergy effectively constitutes a barrier to cooperation because it is rarely achieved in practice. Here we show that punishment reduces this barrier, and enables a transition from defecting towards cooperative behavior at synergy levels that could not support cooperation in the absence of punishment. We use an agent-based evolutionary simulation in which the agents’ decisions to cooperate and to punish are encoded by genes that evolve via Darwinian evolution. We observe that punishment is beneficial for the evolutionary transition from defection to cooperation, but that once cooperation is established the punishment gene becomes unnecessary and drifts neutrally. Thus, punishment is absent in populations that defect and random in populations that cooperate, but is crucial to catalyze the transition between those regimes, and leads to history-dependent effects. We conclude that punishment can be maintained as a low-cost guarantor of cooperation as long as intermittent eruptions of defection maintain the functionality of the punishing pathways, an equilibrium reminiscent of the establishment of global peace via a policy of nuclear deterrent.

“Tragedy of the commons” is the name given to a social dilemma [1] that occurs when a number of individuals maximize their self-interest by exploiting a public good, and by doing so harm their (and others’) own long-term interest. This is but one dilemma [2] that can be described within the framework of Evolutionary Game Theory (EGT) [3–7]. While the tragedy of the commons is important in social science and politics (overfishing, and the destruction of the environment in general come to mind), it also plays an important role in biology: both the evolution of virulence [8] and the manipulation of a host by a group of parasites [9] can be viewed as a dilemma of the public goods type.

The public goods game is a standard of experimental economics [10–12], where players possess a number of tokens that they can contribute to a common pool (the “investment” into the public good). The total contributed by the players is multiplied by a “synergy factor”, and this amount is then equally distributed to the players in the pool, irrespective of whether they have contributed or not. A group of players fares best if all the players contribute so as to take maximum advantage of the synergy, but this

behavior is vulnerable to “free-riders” that share in the pool but do not contribute themselves. Indeed, as can easily be shown, the rational Nash equilibrium of the game is to not to pay in, because this strategy dominates all others regardless of their play.

It has been shown that *punishment* is an effective way to counteract defectors [13–23]. Because punishment involves an additional cost to the co-operators that already invest into the public good [24–26], these cooperators (termed “moralists” by Helbing et al. [23]) are themselves vulnerable to the invasion of non-punishing cooperators called “secondary free-riders”. As a consequence, we might expect that moralists ultimately become extinct, either because they were outcompeted by defectors, or by cooperating free-riders who benefit from the punishment without the associated cost. On the other hand, if moralists were ultimately successful in eliminating defectors, then there would effectively be no difference between cooperators that punish and those that do not, as no punishment ever takes place. Thus, in no case would punishment ever become the dominant strategy, but would only play a role at the boundary between defection and cooperation.

It was recently shown that, instead, in the spatial version of the public goods game, moralists can win direct competitions [23] if the environmental conditions are favorable, namely if the cost and effect of punishment favors moralists over defectors. Spatial games, where the offspring of successful strategies are placed near the parent—and where as a consequence strategies are more likely to play against kin strategies—give rise to spatial reciprocity [19], which appears to be the advantage that moralists need to gain superiority.

According to References [23, 27, 28], punishing cooperators do not fare so well, in contrast, in well-mixed populations. There, punishing cooperators appear to lose the fight against the cooperators that do not punish, and that catch a “free ride”, as it were, on the costly punishment meted out by their moralist peers. As a consequence, defectors can spread. This is a surprising conclusion if it is meant to be unqualified, because at the very least it should be clear that, as long as the synergy gain is large enough, cooperators will be favored no matter whether the dynamics are spatial or well-mixed. Indeed, for a game with 5 players in the group for example (the case studied here and elsewhere), a five-fold synergy implies that the total payoff for a defector is the same as for a cooperator independent of the group he is in, and if the synergy is higher then that, not paying in is in fact detrimental to the individual.

We believe that the solution to both conundrums—the survival of the moralists in the spatial game and the ineffectiveness of punishment in the well-mixed game—can be solved if punishment is not a binary choice (you are either a punisher or not), but is instead a stochastic decision where the probability to punish is shaped by the evolutionary process. Here, we show that if punishment is stochastic, spatial reciprocity is in fact not a necessary condition for the evolution of cooperation via punishment and the dominance of moralists. If stochastic strategies can evolve via Darwinian dynamics in a framework where decisions are encoded within genes that adapt to their environment, we can find conditions where cooperation evolves even without punishment, but absent those, punishment can promote the evolution of cooperation (as long as punishment is effective and cheap) in well-mixed populations.

In previous work, we have investigated the evolution of stochastic strategies in the iterated Prisoner’s Dilemma where players’ decisions are conditional on their previous behavior [29], and found that cooperation is favored as long as the communication channel between players was reliable enough. In a sense, the public goods game is a multi-player Prisoner’s Dilemma so we should expect similar dynamics, except that players in the public goods game do not remember previous plays. Thus, cooperation has to be ensured by different means, for example by punishment. Still, many of the characteristics that we found in the stochastic implementation with a genetic basis we will encounter here too: strategies defined by genes encoding decision probabilities evolve towards a fixed point that is optimal given the selective pressures and environmental conditions. However, the selective pressures are determined by the population: if defectors are absent, for example, genes encoding probabilities that are only “expressed” if defectors are present drift neutrally. Thus, we do not expect that punishing cooperators are maintained after defectors have been driven to extinction in this scenario. When punishment is meaningless, it becomes random.

However, we will see that punishment is critical in the transition from defection to cooperation, playing the role of a catalyst.

Results

Evolutionary trajectories and fixed points

We evolve stochastic strategies playing the public goods game with punishment in a well-mixed population, as described in Methods. Agents possess two genes: one (from now on called the “C gene”) defining the probability to cooperate p_C , and a gene that determines the probability p_P to punish (from now on called the “P gene”). As the strategies adapt to the environmental conditions (specified by the parameters that define the game, as well as the spatial properties, the mutation rate, and the replacement rate), the probabilities change from their initial values $(p_C, p_P) = (0.5, 0.5)$ towards the selected “fixed point” strategy. In order to visualize the evolutionary trajectory of a population, we reconstruct the evolutionary line of descent of an experiment (LOD, see Methods), which tells the story of that adaptation, mutation by mutation. While the LOD in each particular run can show probabilities varying wildly, averaging many such LODs can tell us about the selective pressures the populations face. In particular, averaging the probabilities on the LODs after they have settled down, can tell us the *fixed point* of evolutionary adaptation [29]. We determine this fixed point by discarding the first 250,000 updates of every run (the transient), along with the last 50,000 (in order to remove the dependence of the LOD on the randomly chosen anchor genotype) and averaging the remaining 200,000 updates. Note that this fixed point is a computational fixed point only: we do not mean to imply that the population’s genotypes all end up on this exact point. Rather, due to the nature of the game and the selective pressures that change as the composition of the population changes, the evolutionary trajectories approach this point and then fluctuate around or near it. Thus, the fixed point reflects the *mean* successful strategy given the conditions of the game.




Figure 1. Evolutionary trajectories for different synergies. Evolution of strategies (p_C, p_P) on the LOD for synergy factors $r = 3$ (black), $r = 4$ (green), and $r = 5$ (red). All trajectories originate at $(0.5, 0.5)$. We show an average of the LOD of 10 runs each. Here, $\beta = 0.8$, $\gamma = 0.2$, and $\mu = 0.02$.

We show in Fig. 1 the average trajectories for three different synergy factors $r = 3, 4$, and 5 , all anchored at the random strategy $(p_C, p_P) = (0.5, 0.5)$ that was used as the seed strategy for every

evolutionary run. We can see that, depending on the synergy (and the values chosen for the cost and effect of punishment), populations evolve towards a cooperating or defecting fixed point, and take different trajectories to get there. For $r = 3$, synergy is too low to lead to cooperation, and the fixed point of that trajectory is $(p_C, p_P) = (0, 0)$, that is, defection. For $r = 4$, however, the population moves toward a fixed point centered around $(p_C, p_P) = (0.7, 0.2)$, that is, players cooperate most of the time. (The location of the endpoint of the trajectory does not depend on the starting point.) Note, however, that the players engage in punishment only sparingly. For $r = 5$, cooperation is almost fully established, while punishment occurs about 40% of the time on average. However, the average trajectory (average over ten independent runs) is misleading here, because at this level of cooperation, the punishment gene has begun to drift. This is due to a substantially weakened selection on the punishment gene if players engage in defection only 5% of the time. An unselected probability p_P is a uniformly distributed random number, with mean $1/2$ and variance $1/12$. As $p_C \rightarrow 1$, the average p_P and its variance approach precisely these numbers.

When mapping the possible parameters β (effectiveness) and γ (cost) of punishment (defined in Methods) each in the range from 0.0 to 1.0 and at low synergy $r = 3.0$, we find that defection is the most prevalent strategy on the LOD (see Figure 2A), as was found previously [22, 23]. When $\gamma = 0$ there is no cost associated with the punishment, which implies that the P gene is not under selection and drifts. Thus, for this value of synergy (and lower), we find that the strategy fixed point is defection without punishment, except for the values $\gamma = 0$, where punishment is random.

As the degree of synergy increases to $r = 3.5$, cooperation starts to appear even in this well-mixed population (see Fig. 2B), while it appears as early as $r = 2$ for sufficiently high β and low γ in the spatial (but deterministic) version of the game, see [22, 23]. For $r = 4$ we find players cooperating ($p_C \approx 0.8$) at high β and low γ which indicates that under conditions where punishment is not very costly or even free, punishment pays off. In addition we notice that the probability to punish increases under the same conditions that allows cooperation (high β and low γ , that is high impact, low cost of punishment), indicating that punishment is indeed used to enforce cooperation (Fig. 2C). The mean punishment probability grows to 0.5, but at the same time the variance shows that this gene is not under selection (as long as $\gamma \neq 0$).

Increasing the synergy level even higher towards $r = 4.5$ shows the emergence of dominance of cooperation ($p_C > 0.5$) for most of the range of punishment cost and effectiveness, see Figure 2D. At the same time the punishment probability reaches 0.5 for a larger range of parameters, but the mean punishment probability on the LOD never exceeds 0.5, implying that full persistent punishment is not stable, and probably not necessary. Note that, in an implementation where decisions are deterministic (such as in the implementation of Helbing et al. [23]), punishment may remain for a long time in the population even though it is not selected anymore. In that case, players that cooperate with and without punishment have exactly the same fitness, and one or the other strategy should only dominate by drifting to fixation neutrally, a process that can take a significant amount of time in large populations such as those studied in Ref. [23].

Critical dynamics and the role of punishment

Previously, a phase transition between cooperative and defective behavior in the public goods game as a function of the synergy r was observed for the spatial version [22, 28, 30] of the game (but not the well-mixed version). We can study the critical point and its dependence on punishment in detail in the well-mixed version of the game, where analytical predictions are available. We show in Fig. 3 the average probability to cooperate (solid line) and to punish (dashed line) as a function of synergy for our default values $\gamma = 0.2$ and $\beta = 0.8$. Cooperation sets in at $r = 4$ and becomes prevalent for synergies just exceeding that.

We will now study how punishment affects the critical point. The average probability of cooperation in Fig. 3 shows the typical behavior of an order parameter as a function of the critical parameter r . It

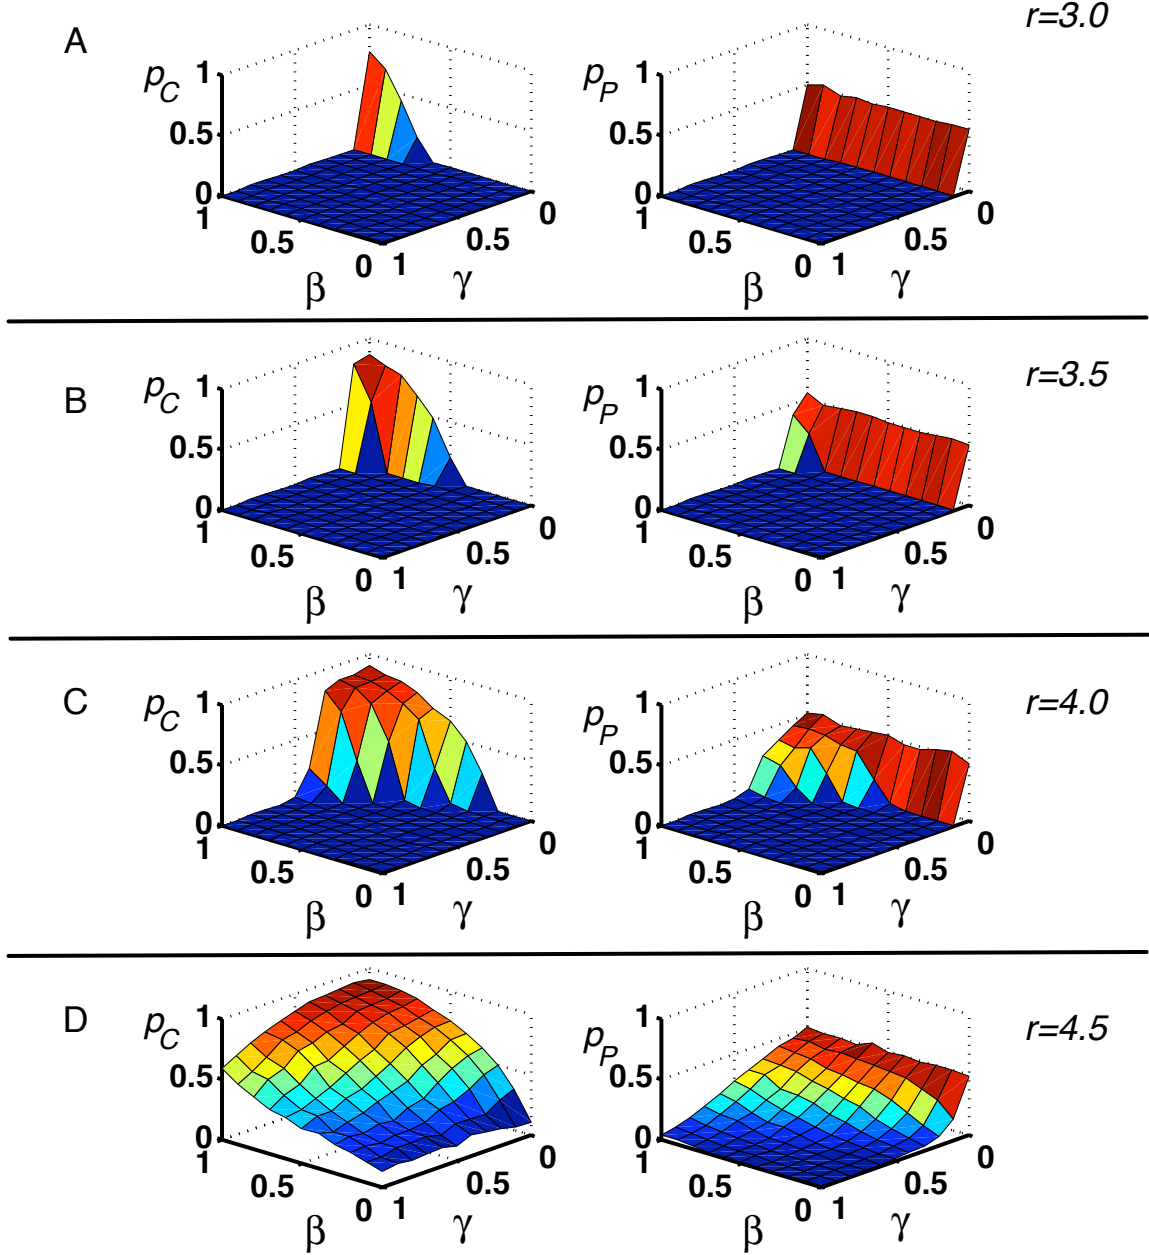


Figure 2. Mean probabilities for cooperation p_C and punishment p_P at the evolutionary fixed point. These graphs show the fixed point as a function of the cost of punishment γ and the effectiveness of punishment β , for different values of the synergy r . Left panel: probability to cooperate p_C , right panel: probability to punish p_P . Note the inversion of the β and γ scales for better visibility. Mutation rate is set to $\mu = 0.02$. **A:** For $r = 3$, cooperation does not evolve except when punishment is free ($\gamma = 0$), and even then only if punishment is very effective (β close to 1). At $\gamma = 0$, the punishment gene is neutral. **B:** For $r = 3.5$ defection is still the predominant strategy except for very low γ and high β . **C:** At $r = 4$, cooperation is fully established for low γ and high β , but not for medium values. **D:** For $r = 4.5$ cooperation is the dominant strategy for all values of the cost γ , and for high effect ($\beta > 0.75$). Note that the average punishment probability p_P never exceeds 0.5 (the value achieved when the gene drifts neutrally).

is instructive to run a control of the experiment where punishment does not exist. If we force $p_P = 0$, cooperation does not set in until $r = 4.5$ (see inset in Fig. 3) and only becomes dominant at $r = 5$. Thus, although punishment is sporadic when it is possible—and drifts when cooperation is established—it is essential to lower the critical barrier for cooperation. The probability distribution of the punishment gene throughout the population (Fig. 4) shows that punishment is never prevalent: it is absent below the critical point, and close to uniform above it. In a sense, punishment catalyzes the transition from defection to cooperation. Note also that the levels of cooperation achieved are significantly higher when punishment exists, even though punishment is only weakly selected for. Apparently, the possibility of punishment alone is sufficient to enforce higher levels of cooperation.




Figure 3. Mean probability of cooperation p_C (solid, left scale) and probability of punishment p_P (dashed, right scale) with adaptive punishment at the evolutionary fixed point of the trajectory, as a function of the synergy r ($\beta = 0.8, \gamma = 0.2, \mu = 0.02$ (100 replicates for each data point)). The probability to cooperate when punishment is forced to zero ($p_P = 0$) is shown in the inset.

We can calculate approximately the point at which cooperation is favored in a mean-field approach that does not take mutation and evolution into account, by writing Eqs. (5-6) in terms of the density of cooperators ρ_C encountered by players in a group. Both naked cooperators and punishing cooperators (moralists) contribute to this density, i.e., $\rho_C = (N_C + N_M)/N$, where N is the total number of players in the group. We can also introduce the mean density of punishers $\rho_P = (N_M + N_I)/N$ encountered by a player. Because the mean density of cooperators and punishers is the *same* for both cooperators and defectors in a well-mixed scenario (but not for spatial play!), we can then write

$$P_C = r \frac{k\rho_C + 1}{k + 1} - 1 \quad (1)$$

and

$$P_D = r \frac{k\rho_C}{k + 1} - \beta\rho_P, \quad (2)$$

and we expect cooperation to be favored if

$$P_C - P_D = \frac{r}{k + 1} - 1 + \beta\rho_P > 0 \quad (3)$$

or

$$r > (k + 1)(1 - \beta\rho_P). \quad (4)$$

This equation implies that the emergence of cooperation depends crucially on the density of punishers.




Figure 4. Histogram of the punishment probability distribution, in a typical equilibrated population, just before the critical point ($r = 4$, red), at the critical point ($r = 4.15$, green), and above r_{crit} ($r = 4.5$, blue).

In fact, the mean-field theory predicts that cooperation in the absence of punishment is favored only at $r = 5$. We see cooperation emerge quite a bit earlier than that in our simulations (see inset in Fig. 3), but crosses $p_C = 0.5$ very close to $r = 5$, as predicted by the mean field theory.

We can test Eq. (4) by finding the critical r at which p_C crosses 0.5 for simulations in which the punishment probability is held fixed, so that $\rho_P \approx p_P$. To find the critical point, we performed 100 simulations each at fixed r with a resolution of $\Delta r = 0.5$ and interpolated data within the steep portion of the transition to find the crossover point. The critical line $r_c = (k + 1)(1 - \beta p_P)$ is indicated in Fig. 5 for $k = 4$ and $\beta = 0.5$ ($r_c = 5 - 4p_P$). The mean field theory reproduces the experimental r_c within errors.




Figure 5. Prediction of critical point at fixed punishment [Eq. (4), solid line] and extrapolated critical point at transition, for simulations in which the probability to punishment was kept fixed and constant. We used $k = 4$, $\beta = 0.5$, and $\gamma = 0.2$. The error bars reflect the finite resolution $\Delta r = 0.5$.

Because of the critical importance of punishers in determining the synergy level at which cooperation

emerges, the public goods game with a genetic basis implies curious dynamics close to the critical point. Below the critical point, defection is a stable strategy, and punishment is absent. Only when cooperation emerges as a possibility, punishment becomes more and more important, leading to a lowering of the critical synergy for cooperation via Eq. (4). Thus, cooperation emerges rapidly and decisively once a critical level has been achieved. Once cooperation is dominant and defectors are all but driven to extinction, punishment becomes irrelevant and the gene begins to drift. As this happens, the fraction of punishers drops, raising the critical synergy. Thus, a drifting punishment gene can lead to the sudden re-emergence of defectors as stable states. Once those have taken over, the reverse dynamics begins to unfold. In other words, we should observe periods of cooperation and defection that follow each other closely when the synergy is near the critical point.

These dynamics are reminiscent of the phenomenon of supercooling and superheating in phase transitions. If we imagine the synergy parameter r as the critical parameter and the mean probability to cooperate as the order parameter, it is possible that when r is slowly increased, the population remains in the defecting phase because a switch to cooperation requires a critical number of cooperators as a “seed”. In such a situation, the defecting phase is unstable to fluctuations. If a critical number of cooperators emerges by chance, punishment immediately becomes effective against defectors, lowers the critical point as implied by Eq. (4), and the population could transition to cooperation very quickly. A hallmark of such bi-stable systems that require nucleation events in order to transition is *hysteresis*, a phenomenon where the state of the system depends on its history. We can test whether hysteresis exists in the public goods game (and whether the strength of this effect depends on the probability to punish), by adiabatically changing the synergy parameter first from low to high (transitioning from defection to cooperation), and then adiabatically back from high to low. While we see evidence of hysteresis even when punishment




Figure 6. Population fraction of cooperators (measured as the density of non-punishing cooperators plus the density of moralists) as a function of synergy r when r is adiabatically changed from low to high values (red), and back from high values to low values (blue). All population fractions are started at 0.5 (either at the high or low end of r). The lines show the average over 100 runs. Standard error is of the size of the fluctuations.

is absent (Fig. 6A), the effect is much more pronounced when punishment is possible (Fig. 6B). The population moves from cooperation to defection at about the expected critical synergy $r_{\text{crit}} \approx 4.15$ as r is decreased, but stays in the defecting phase much beyond the critical point as r is increased.

Discussion

We studied Darwinian evolution of stochastic strategies in the public goods game for well-mixed populations, using genes that encode the probabilities for cooperation and punishment. It is known that

punishment can drive the evolution of cooperation above a critical synergy level as long as there is a spatial structure in the environment [22, 23]. It was also previously believed that in well-mixed populations cooperation can only become successful if additional factors like reputation [19] or the potential for abstaining from the public good [30, 31] are influencing the evolution. Here we show that cooperation readily emerges in a well-mixed environment above a critical level of synergy. This critical level is influenced by a number of factors: the rate of punishment because punishment favors cooperating groups, but also mutation rate, because mutations can create cooperating groups by chance, and encourage a minimal punishment rate in order to keep mutated defectors at bay. Finally, spatial structure also affects the critical point as is well known [22, 28, 30], because a single cooperator can nucleate a transition simply because offspring cooperators are placed next to it, giving rise to a “bubble” of cooperators of sufficient size.

We conclude that in well-mixed populations cooperation can emerge if the synergy outweighs the defectors’ reward. If the mutation rate is low enough, the dearth of defectors in the cooperating phase makes punishment obsolete, that is, the selective pressure to punish disappears. Naturally, once this has occurred defectors can again gain a foothold, and the balance of power between cooperators and defectors could shift. Such a shift, however, reinstates the selective pressure to punish, leading to a re-emergence of moralists that can drive defectors out once more. Thus, for synergy factors near the critical point, we can expect oscillations between cooperators and defectors, and no strategy is ever stable.

We have not studied here the possibility of “anti-social” punishment [32], where non-cooperating defectors can punish cooperators, but we do not expect this possibility to change the overall picture. Indeed, in simulations in which defection was not punished but instead rewarded (a negative punishment), this only served to reinforce the defecting phase. A transition to the cooperative phase still takes place at sufficiently high synergy. Phase transitions between cooperative and defection phases have also been observed in a spatial version of the public goods game where costly rewards are given for cooperation, rather than the costly punishment for defectors [33]. It would be interesting to study this game within the context of evolving stochastic strategies.

It is difficult to evade the analogy between punishment as a catalyzing agent of cooperation (while punishment is in fact rarely used), to the politics of a nuclear deterrent and mutually assured destruction, where the threat of severe punishment alone is sufficient to ensure long periods of peace between superpowers. Previously, the game of “chicken” from the EGT literature was used to describe the politics of deterrence [34], but in that game defection affected only the players not an entire community, and the punishment for uncooperative behavior was the action of defection itself. In the public goods game with punishment the punitive action is a reaction to defection, and its threat alone appears to be sufficient to realize peaceful coexistence for prolonged periods of time.

Methods

The public goods game emulates strategic decision making by groups, in which individual must select between different decisions that affect the group as a whole. Each individual in a group of $k + 1$ players ($k = 4$ in the present implementation) can decide to cooperate by making a contribution of 1 unit to the public good, while defecting individuals do not contribute. We encode this choice into a genetic locus as a probability p_C , which can be thought of as the outcome of a network of genes that encode this decision. When mutating strategies, instead of mutating the individual genes that make up the decision pathway, we simply replace the parental probability p_C by a uniformly drawn random number in the offspring.

The sum of all contributions from cooperating players is multiplied by r (the synergy factor) and divided among all players. In addition, each player has the option to punish players who do not contribute. This decision is encoded into a different genetic locus with an independent probability p_P . Following Helbing et al. [23], those players that defect suffer a fine β/k levied by each punisher in the group, which costs each punisher a penalty of γ/k . At each update, every player engages in a game with all its assigned

opponents. The number of cooperators N_C , defectors N_D , moralists N_M and immoralists (players who defect but also punish [23]) N_I is computed, and the payoff is assigned as follows: A cooperator receives

$$P_C = r \frac{(N_C + N_M + 1)}{k + 1} - 1, \quad (5)$$

while a defector takes away

$$P_D = r \frac{(N_C + N_M)}{k + 1} - \beta \frac{(N_M + N_I)}{k}. \quad (6)$$

Moralists receive

$$P_M = P_C - \gamma \frac{(N_D + N_I)}{k}, \quad (7)$$

while immoralists earn

$$P_I = P_D - \gamma \frac{(N_D + N_I)}{k}. \quad (8)$$

The population consists of 1,024 individuals who each have four assigned opponents. Since all opponents are also players, each individual plays five games per update. The choices of each individual are determined by their probabilities to cooperate p_C and to punish p_P . After each round, 2 percent of the population is replaced using a Moran-process [35] in a well-mixed fashion, that is, the identity of the players in the group is unrelated to their ancestry so that, effectively, the members of a particular playing group are randomly selected from the population. We verified that the probability for a player to encounter cooperators is independent of whether that player is a cooperator or a defector, as is required for well-mixed populations [36]. Players that are not replaced are allowed to accumulate their score, which is used to calculate the probability that this player's strategy will be chosen to replicate and fill the spot of a player that was removed in the Moran process. While the spatial version of the game shows somewhat different dynamics than studied here, we study the well-mixed version because it is amenable to theoretical prediction (see below). In fact, cooperation is harder to achieve in well-mixed populations, so most of our conclusions translate to the spatial version but with a lower synergy threshold.

The two genes of every individual mutate with a probability μ when replicated. As mentioned earlier, mutating a probability replaces the probability with a uniformly distributed random number. After 500,000 updates, the line of descent (LOD) of the population is reconstructed [37, 38], by picking a random organism of the final population and following its ancestry all the way back to the starting organism, which has $p_C = 0.5$ and $p_P = 0.5$. Because there is only one species in these populations, the LODs of the population coalesce to a single LOD (which is why it is sufficient to pick a random genotype for following the LOD).

Acknowledgements

This work was supported in part by the National Science Foundation's Frontier in Integrative Biological Research Grant No. FIBR-0527023 and NSF's BEACON Center for the Study of Evolution in Action, under Contract No. DBI-0939454. We wish to acknowledge the support of the Michigan State University High Performance Computing Center and the Institute for Cyber Enabled Research.

References

1. Hardin, G. The tragedy of the commons. *Science* **162**, 1243–1248 (1968).

2. Frank, S. A. *Foundations of Social Evolution* (Princeton University Press, 2006).
3. Smith, J. M. *Evolution and the Theory of Games* (Cambridge University Press, Cambridge, UK, 1982).
4. Axelrod, R. *The Evolution of Cooperation* (Basic Books, New York, NY, 1984).
5. Dugatkin, L. A. *Cooperation Among Animals: An Evolutionary Perspective* (Princeton University Press, Princeton, NJ, 1997).
6. Hofbauer, J. & Sigmund, K. *Evolutionary Games and Population Dynamics* (Cambridge University Press, Cambridge, UK, 1998).
7. Nowak, M. *Evolutionary Dynamics* (Harvard University Press, Cambridge, MA, 2006).
8. Frank, S. Models of parasite virulence. *Quarterly Review of Biology* **71**, 37–78 (1996).
9. Brown, S. Cooperation and conflict in host-manipulating parasites. *Proceedings of the Royal Society of London Series B-Biological Sciences* **266**, 1899–1904 (1999).
10. Olson, M. *The logic of collective action: Public goods and the theory of groups* (Harvard University Press, Cambridge, MA, 1971).
11. Davis, D. D. & Holt, C. A. *Experimental Economics* (Princeton University Press, Princeton, N.J., 1993).
12. Ledyard, J. Public goods: A survey of experimental research. In Kagel, J. H. & Roth, A. E. (eds.) *Handbook of experimental economics*, 111–194 (Princeton University Press, Princeton, N.J., 1995).
13. Fehr, E. & Gächter, S. Altruistic punishment in humans. *Nature* **415**, 137–140 (2002).
14. Fehr, E. & Fischbacher, U. The nature of human altruism. *Nature* **425**, 785–791 (2003).
15. Hammerstein, P. (ed.). *Genetic and Cultural Evolution of Cooperation* (MIT Press, Cambridge, MA, 2003).
16. Nakamaru, M. & Iwasa, Y. The coevolution of altruism and punishment: role of the selfish punisher. *J Theor Biol* **240**, 475–88 (2006).
17. Camerer, C. F. & Fehr, E. When does “economic man” dominate social behavior? *Science* **311**, 47–52 (2006).
18. Gürer, O., Irlenbusch, B. & Rockenbach, B. The competitive advantage of sanctioning institutions. *Science* **312**, 108–11 (2006).
19. Sigmund, K., Hauert, C. & Nowak, M. A. Reward and punishment. *Proc Natl Acad Sci U S A* **98**, 10757–62 (2001).
20. Henrich, J. & Boyd, R. Why people punish defectors. weak conformist transmission can stabilize costly enforcement of norms in cooperative dilemmas. *J Theor Biol* **208**, 79–89 (2001).
21. Boyd, R., Gintis, H., Bowles, S. & Richerson, P. J. The evolution of altruistic punishment. *Proc Natl Acad Sci U S A* **100**, 3531–5 (2003).
22. Brandt, H., Hauert, C. & Sigmund, K. Punishment and reputation in spatial public goods games. *Proc Biol Sci* **270**, 1099–104 (2003).

23. Helbing, D., Szolnoki, A., Perc, M. & Szabó, G. Evolutionary establishment of moral and double moral standards through spatial interactions. *PLoS Comput Biol* **6**, e1000758 (2010).
24. Yamagishi, T. The provision of a sanctioning system as a public good. *Journal of Personality and Social Psychology* **51**, 110–116 (1986).
25. Fehr, E. Human behaviour: Don't lose your reputation. *Nature* **432**, 449–450 (2004).
26. Colman, A. M. The puzzle of cooperation. *Nature* **440**, 744–745 (2006).
27. Helbing, D., Szolnoki, A., Perc, M. & Szabo, G. Defector-accelerated cooperativeness and punishment in public goods games with mutations. *Physical Review E* **81**, 057104 (2010).
28. Helbing, D., Szolnoki, A., Perc, M. & Szabo, G. Punish, but not too hard: how costly punishment spreads in the spatial public goods game. *New Journal of Physics* **12**, 083005 (2010).
29. Iliopoulos, D., Hintze, A. & Adami, C. Critical dynamics in the evolution of stochastic strategies for the iterated Prisoner's Dilemma. *PLoS Computational Biology* **7**, e1000948 (2010).
30. Szabo, G. & Hauert, C. Phase transitions and volunteering in spatial public goods games. *Physical Review Letters* **89**, 118101 (2002).
31. Hauert, C., Traulsen, A., Brandt, H., Nowak, M. A. & Sigmund, K. Via freedom to coercion: the emergence of costly punishment. *Science* **316**, 1905–7 (2007).
32. Rand, D. G. & Nowak, M. A. The evolution of antisocial punishment in optional public goods games. *Nat Commun* **2**, 434 (2011).
33. Szolnoki, A. & Perc, M. Evolutionary advantages of adaptive rewarding. *New Journal of Physics* **14**, 093016 (2012).
34. Rapoport, A. & Chammah, A. N. The game of chicken. *The American Behavioral Scientist* **10**, 10–28 (1966).
35. Moran, P. A. P. *The Statistical Processes of Evolutionary Theory* (Clarendon Press, Oxford, 1962).
36. Fletcher, J. A. & Doebeli, M. A simple and general explanation for the evolution of altruism. *Proc Biol Sci* **276**, 13–9 (2009).
37. Lenski, R. E., Ofria, C., Pennock, R. T. & Adami, C. The evolutionary origin of complex features. *Nature* **423**, 139–144 (2003).
38. Ostman, B., Hintze, A. & Adami, C. Impact of epistasis and pleiotropy on evolutionary adaptation. *Proc Biol Sci* **279**, 247–56 (2012).